

Statement of Teaching Philosophy

Andrey Makhanov

Ph.D. Candidate, Business Administration — Operations Management and Data Analytics

UNC Kenan-Flagler Business School

Andrey_Makhanov@kenan-flagler.unc.edu | andreymakhanov.com

Overview

I teach business analytics, operations, and information systems, and my goal is to prepare students to make sound, evidence-based decisions in organizations where data, algorithms, and digital risk have become inseparable from strategy. My approach grows out of a simple conviction that I share with the research side of my work: models matter only insofar as real people use them, and people exercise judgment. A regression, an optimization model, or a security control framework does not make a decision; a manager does, under time pressure, with incomplete information, and with incentives pulling in several directions at once. Good teaching in a business school therefore has to do two things at the same time. It has to build genuine technical fluency, and it has to build the judgment that determines whether that fluency is used well. Everything else in my classroom follows from that pairing.

Three commitments organize how I put this into practice: anchoring every technical concept to a decision someone actually has to make, building hands-on capability with the tools students will use the week after graduation, and holding high expectations for students who arrive with very different quantitative preparation. These commitments were shaped by teaching Business Analytics as a lecturer at UNC Kenan-Flagler, by serving repeatedly as a teaching assistant across analytics, prescriptive modeling, data tools, and operations courses for undergraduate, graduate, and professional audiences, and by more than a decade of applied work in machine learning, software engineering, and security before and alongside my doctoral training.

Anchoring Concepts to Decisions

Business students learn a method best when they first feel the absence of it. When I taught Business Analytics, I built each module backward from a managerial question—How should we price this? Which supplier should we choose? Is this investment in reliability worth what it costs?—and introduced the quantitative machinery only once the question had become uncomfortable to answer by intuition alone. Regression, optimization, simulation, and classification then arrive as answers to a problem the students already own rather than as techniques to be memorized and reproduced on an exam. What I ask of them at the end of a session is not only that they can explain how a model works, but that they can say when to trust it, how to communicate its output to someone who will never see the code, and what its limitations imply for the decision at hand.

That last point is where I spend the most class time, because it is where students are weakest and where practice is least forgiving. A model that is technically correct and rhetorically unusable has no effect on an organization, and a model whose assumptions have quietly stopped holding is worse than no model at all. So I ask students to defend recommendations out loud, to state what would change their mind, and to identify the one assumption whose failure would most damage their conclusion. The same logic scales up from an analytics exercise to a board-level question about risk exposure or technology investment: frame the decision first, build the analytical apparatus second,

interpret the two together third.

Building Capability with Modern Tools

Students in analytics and information systems have to build things with their own hands. In my courses and teaching-assistant sections I work directly in Python, R, SQL, Excel, and standard visualization and modeling environments, and I draw assignments from realistic data rather than from sanitized examples that behave. Messy data is not an obstacle to the pedagogy; it is much of the pedagogy. When a join silently drops a third of the rows, or a variable turns out to be recorded differently in two systems, students learn something about analytical practice that no clean dataset can teach them.

My own technical background informs the exercises I design. Work spanning machine learning research, Kubernetes-based reproducible research computing at UNC's Odum Institute, production software and identity-verification systems in industry, and graduate coursework in network security, malware analysis, applied cryptography, secure systems, and technology policy has given me a concrete sense of where analytical and technical work actually breaks. I try to put students in front of those failure points deliberately and early, at a scale they can recover from, so that the experience becomes instructive rather than discouraging. Where a course calls for it, I add short applied artifacts: reading real disclosures and filings, mapping organizational scenarios onto standard control and governance frameworks, quantifying exposure with loss models, working through compliance cases under existing privacy regimes, and running structured tabletop exercises. Small technical labs—log analysis, reconnaissance concepts, adversarial machine learning demonstrations—serve the same purpose of letting students see the mechanism beneath a decision they will one day make from a distance.

Analytics, AI, and Risk as Decision Problems

A distinctive feature of my teaching is that I treat technical subjects as decision problems involving uncertainty, incentives, human behavior, and organizational design. In a business school, the students who will make the most consequential decisions about analytics, AI, and technology risk are rarely the ones who will write the code. They are the managers and executives who decide how much to invest, which vendors to trust, how much automation to accept, how to disclose an incident, and how to balance control against usability, innovation, and cost. Teaching them to read a model output is necessary but insufficient; they need frameworks for acting when the output is ambiguous.

My courses therefore give sustained attention to reasoning under uncertainty through expected loss, decision trees, and scenario analysis; to governance structures and escalation paths; to the regulatory and disclosure environment surrounding data and technology; to responsible use of AI systems, including model risk and data leakage; and to the human dimensions of organizational risk, from behavioral compliance to insider threat. These topics reward the same treatment as any operations problem. The technical constraint is real, but the outcome usually depends on how people interpret the information and the incentives in front of them.

Responsible Use of AI in Business Education

Generative AI has changed what it means to teach quantitative business courses, and I have found it more productive to integrate these tools explicitly than to police them. Students in my classes learn to prompt, critique, and verify AI-generated code and analysis; to recognize confident errors and data-leakage risks; to document where and how they used assistance; and to judge when a task can be delegated to a model and when it requires human judgment. The exercise is deliberately

uncomfortable: I ask students to find the flaw in output that looks polished, because that is precisely the skill their employers will need from them. This mirrors professional reality and prepares them to help set responsible AI policy rather than merely comply with it.

Teaching Across Formats

Business education now happens in several modalities at once, and I have prepared for each. My own graduate training in analytics, computer science, and cybersecurity was delivered primarily online and at scale, which gave me an unusually direct student-side view of what online technical education does well and where it fails—what a discussion board can and cannot accomplish, how quickly a poorly scoped asynchronous assignment becomes unrecoverable, and how much scaffolding remote students need before they will ask for help. I have taught in person at UNC Kenan-Flagler and supported technology-enabled instruction as a teaching assistant. I design activities that translate across formats: short pre-class material paired with problem-based synchronous sessions, asynchronous prompts anchored to current events and cases, team projects built on shared cloud notebooks and version control, and rubric-based feedback that holds up as cohort size grows.

Assessment and Continuous Improvement

Assessment should measure what actually matters: whether a student can frame a problem, choose a defensible method, execute it competently, and communicate the result to someone who has to act on it. My courses combine short skill checks that establish fluency, integrative cases and projects that build judgment, and brief reflective components that develop students' sense of when and why to use a given method. I give detailed individual feedback, and I revise materials each term based on evaluations, observed learning outcomes, and conversations with students during office hours. I treat my own teaching the way I treat my research—as an evolving practice that should answer to evidence.

Inclusive and Rigorous Classrooms

Business students arrive with very different quantitative preparation, professional experience, and cultural background, and a course designed for the median student serves almost no one well. My classrooms are built so that the student who has never written a line of code and the student with a computer science degree both find real challenge and real support: scaffolded assignments with optional stretch problems, structured rather than incidental peer collaboration, worked examples that expose reasoning rather than only results, and office hours that are clearly signposted and genuinely used. High expectations and substantial support are not in tension; separating them is what makes rigor exclusionary. My service and leadership roles—as Ph.D. Student Body Co-President and Senator in UNC's Graduate and Professional Student Government, as MBASA Vice President of Communications and Technologies, and as a Dean's Fellows project leader—have reinforced how much a student's willingness to ask for help depends on whether the environment has signaled that asking is normal.

Courses and Contributions

I am prepared to teach business analytics, statistics and data-driven decision making, operations and supply chain management, business programming and data management, and courses in technology and security management, at the undergraduate, M.B.A., specialized master's, and executive levels. I would also welcome the chance to develop electives at the intersection of AI governance, analytics, and technology strategy, and to advise applied master's projects and doctoral students working on

data-driven questions in operations and information systems. Beyond the classroom, I expect to contribute where my technical and professional background is useful: curriculum development for analytics sequences, tooling and infrastructure decisions for hands-on courses, and connections to practitioners whose problems make good teaching material.

Closing

My aim is to graduate students who reason quantitatively, act ethically, communicate clearly, and lead confidently in organizations where analytics, AI, and digital risk sit close to strategy. Discretion is the thread running through both my research and my teaching: systems perform well when the people inside them are given good information, sound frameworks, and the judgment to use both. I would be glad to bring that view to your school and to undergraduate, M.B.A., specialized master's, and executive programs.